

The Internet Protocol Journal

December 2008

Volume 11, Number 4

A Quarterly Technical Publication for
Internet and Intranet Professionals

In This Issue

From the Editor	1
Wi-Fi, Bluetooth and WiMAX.....	2
The End of Eternity	18
Remembering Jon.....	29
Letters to the Editor.....	33
Book Reviews	36
Fragments	41
Call for Papers.....	43

You can download IPJ
back issues and find
subscription information at:
www.cisco.com/ipj

ISSN 1944-1134

FROM THE EDITOR

Response to our use of a new printing paper has been very positive, so we will continue to use the uncoated and recycled Exact® paper introduced with our September 2008 issue. We are still interested in hearing your feedback on the paper, as well as any other aspect of this journal. Send your comments to: ipj@cisco.com

The last decade has seen many developments in the area of *wireless* networking technologies. Wireless Internet access is now available in thousands of locations ranging from private homes to hotels, trains, airplanes, ships at sea, and even entire cities. Wireless systems, specifically Bluetooth, are also used for short-range device connectivity such as between a mobile phone and a headset, while WiMAX systems are being deployed for larger area coverage. In our first article, T. Sridhar gives an overview of Wi-Fi, Bluetooth, and WiMAX.

As stated in our previous issue, the topic of IP Version 4 address exhaustion and migration to IP Version 6 is being debated in many Internet-related organizations, including the IETF, *Internet Corporation for Assigned Names and Numbers* (ICANN), and the *Regional Internet Registries* (RIRs). In our last issue, Geoff Huston outlined the history of IPv4 address depletion. This time we bring you the first in a two-part series of articles entitled “The End of Eternity.” The article is by Niall Murphy and David Wilson. Part Two will follow in our March 2009 issue. As you will see from our “Letters to the Editor,” views on the right way to tackle the address exhaustion and protocol migration challenge abound, and I predict we will carry yet more articles on this topic in future issues.

Just over 10 years ago, Jonathan B. Postel, Internet pioneer and a key player in many core Internet activities, passed away. In this issue we bring you a remembrance article written by another Internet pioneer, Vint Cerf. In connection with this anniversary, special events were held in Minneapolis in conjunction with the 73rd meeting of the IETF. The *Jonathan B. Postel Service Award* for 2008 was awarded to EsLaRed of Venezuela by a committee of former award winners. You will find more information about the award in our “Fragments” section on page 42.

Remember to let us know if your mailing address changes and to visit our online companion, *The Internet Protocol Forum*, where you will find additional articles and other material: <http://ipjforum.org>

—Ole J. Jacobsen, Editor and Publisher
ole@cisco.com

Wi-Fi, Bluetooth and WiMAX—Technology and Implementation

by T. Sridhar, Flextronics

Wireless networks can be classified broadly as *Wireless Personal-Area Networks* (WPAN), *Wireless LANs* (WLANs), and *Wireless Wide-Area Networks* (WWANs). WPANs operate in the range of a few feet, whereas WLANs operate in the range of a few hundred feet and WWANs beyond that. In fact, wireless WANs can operate in a wide range—a metropolitan area, cellular hierarchy, or even on intercity links through microwave relays.

This article examines wireless technologies for the WLAN, WPAN, and WWAN areas, with specific focus on the IEEE 802.11 WLAN (often known as Wi-Fi®), *Bluetooth* (BT) in the WPAN, and WiMAX for WWAN as representative technologies. It discusses key aspects of the technology—medium access and connectivity to the wired network—and concludes by listing some common (mis)perceptions about wireless technology.

WLANs

The *Institute of Electrical and Electronic Engineers* (IEEE) defined three major WLAN types in 802.11–802.11 b and g, which operate in the 2.4-GHz frequency band, and 802.11a, which operates in the 5-GHz band. The 2.4- and 5-GHz bands used here are in the license-free part of the electromagnetic spectrum, and portions are designated for use in *Industrial, Scientific, and Medical* (ISM) applications—so these portions are often called ISM bands. More recently, a high-speed 802.11 WLAN has been proposed—the 802.11n WLAN, which operates in both the 2.4- and 5-GHz bands.

The 2.4-GHz frequency band used for 802.11 is the band between 2.4 and 2.485 GHz for a total bandwidth of 85 MHz, with 3 separate nonoverlapping 20-MHz channels. In the 5-GHz band, there are a total of 12 channels in 3 separate subbands—5.15 to 5.25 GHz (100 MHz), 5.25 to 5.35 GHz (100 MHz), and 5.725 to 5.825 GHz (100 MHz).

The more common mode of operation in 802.11 is the *infrastructure* mode, where the stations communicate with other wireless stations and wired networks (Ethernet typically) through an *access point*. The other mode is the *ad-hoc* mode, where the stations can communicate directly with each other without the need for an access point; we will not discuss this mode in this article. The access point bridges traffic between wireless stations through a lookup of the destination address in the 802.11 frame (see Figure 1a).

The *Media Access Control* (MAC) header of 802.11 has four addresses. Depending upon the value of a *FromDS* (from access point), or a *ToDS* (to access point) bits in the header (see Figure 1b), the addresses have different connotations. The first two addresses are for the receiver and transmitter, respectively.

Figure 1a: WLAN Network with Ethernet Connectivity

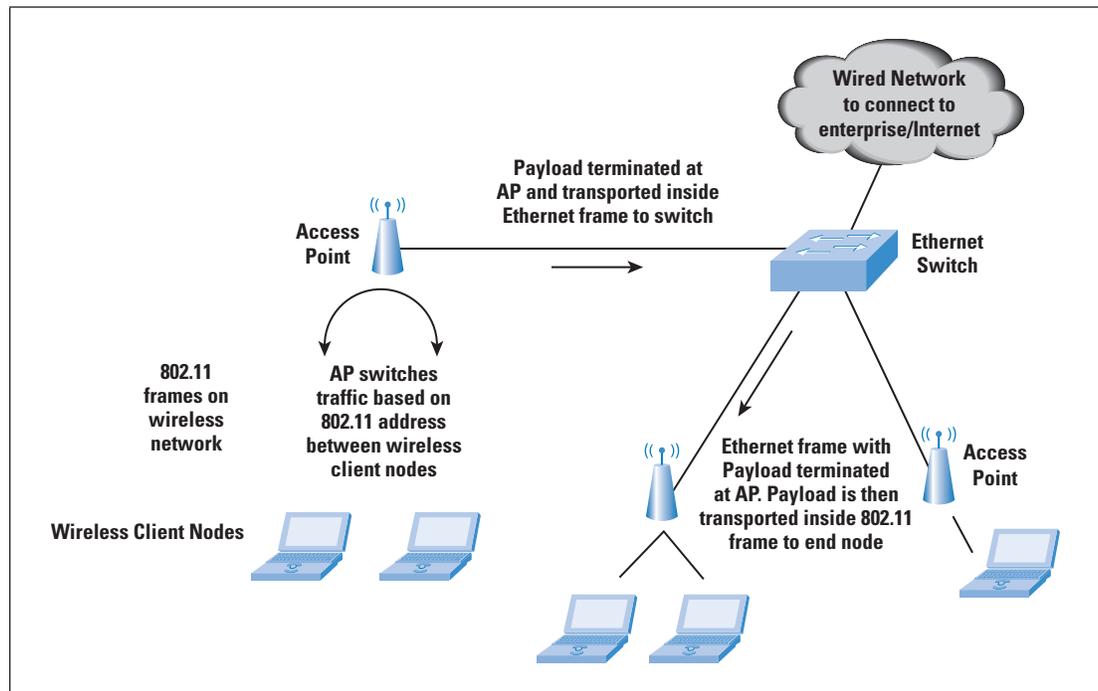
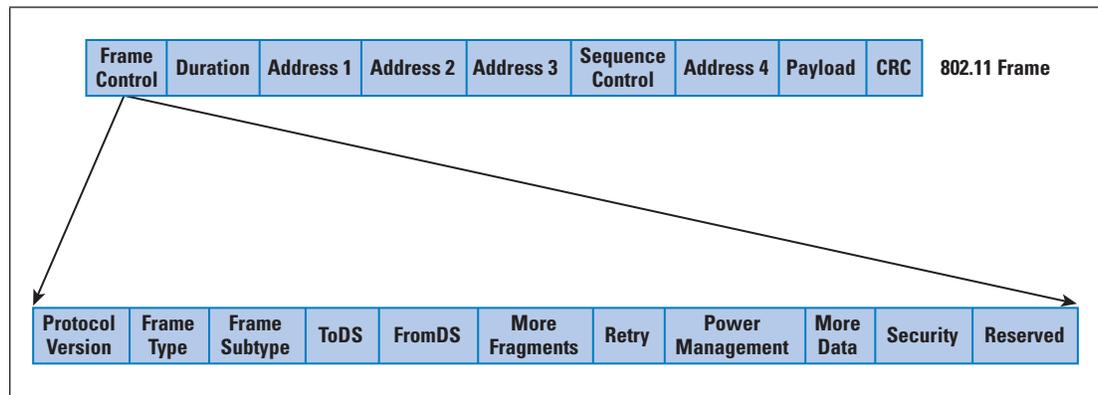


Figure 1b: 802.11 Frame Format



Address 4 is not used except when both FromDS and ToDS are set to 1—it is for a special mode of communication for access point-to-access point traffic, whence addresses 3 and 4 refer to the source- and destination-station MAC addresses, respectively, whereas addresses 1 and 2 refer to the access point addresses (that is, the transmitter and receiver on this inter-access point channel). When FromDS is set to 1, address 1 is the destination-station MAC address, address 2 is the access point address, and address 3 is the source-station MAC address. When ToDS is set to 1, address 1 is the access point MAC address, address 2 is the transmitting-station MAC address, and address 3 is the destination-station MAC address.

Although earlier versions of 802.11 LANs used *Frequency Hopping Spread Spectrum* (FHSS), 802.11b typically uses *Direct Sequence Spread Spectrum* (DSSS) for 1-, 2-, 5.5-, and 11-Mbps speeds. Both schemes involve transmission of a narrowband signal over a wider frequency range to mitigate the possibility of interference at any one frequency. The nodes and access points typically transmit at the highest data rate possible based on the current signal-to-noise ratio.

At the MAC level, 802.11 LANs involve the use of *Carrier Sense Multiple Access/Collision Avoidance* (CSMA/CA). Stations back off if they detect that another station is transmitting on that channel. The station then waits for a random period after the end of the transmission before it attempts to transmit on that channel. In addition, control frames such as *Request to Send* (RTS) and *Clear to Send* (CTS) are used to facilitate the actual data transfer. The CTS control frame has the duration for which the transmitting node is allowed to transmit. Other stations sense this frame and back off for at least the specified duration before sensing the radio link again.

When the access points are connected through a LAN, the entire system is known as a *Distribution System*. The access points perform an integration function—that is, bridging between wired and wireless LANs. In this scenario, (see Figure 1a) the wireless control and data frames are terminated at the access point or tunneled from the access point to a centralized controller over Ethernet. When terminated at the access point, the payload is transmitted from the access point to the network over Ethernet. This transmission is done in the following manner:

The source and destination addresses are set to the station and access point addresses, respectively. At the access point, the payload is stripped from the 802.11 data frame and sent as part of an Ethernet packet either as a broadcast packet or to a specific destination. If the packet sizes (when reassembled) are larger than the Ethernet frame size, they are discarded. In the reverse direction, the Ethernet frame can be directly encapsulated into an 802.11 frame for transmission from the access point to the end node. At the WLAN end node, the complete Ethernet frame shows up at the driver level as though it were a frame received on a pseudo Ethernet interface.

The most common 802.11b WLAN speed is 11 Mbps. However, based on the interframe spacing, preamble, header encapsulation, and acknowledgements for frames required, the actual throughput for user data would be about 50 percent of the actual speeds. This throughput of 50 percent of actual link speed is a common theme on 802.11g and 802.11a also.

Stations connect to the access point through a scanning process. Scanning can be passive or active. In the passive mode, the station searches for access points to find the best access point signal (which contains the *Service Set Identifier* [SSID], data rates, and so on).

The access point frame that the stations look for is a management frame known as the *beacon frame*. In the active mode, the station initiates the process by broadcasting a probe frame. All access points that receive the probe send back a probe response, helping the station build up the list of available access points. The sequence of a station “connecting” to an access point involves two steps. The first is *authentication*, where the station sends an authentication request frame to the access point. Depending upon the authentication through 802.1X or internal configuration, the access point can accept or reject the request with an authentication response. The second step is *association*, which is required to determine the data rates supported between the access point and the station. At the end of the association phase, the station is allowed to transmit and receive data frames.

Power Concerns in 802.11

Although it is not a part of the standard, the access points might adjust their transmitting power based on the environment they are in (they do have maximum limits based on regional restrictions). If they do not perform this adjustment, all the stations might connect to the access point with the highest transmitting power, even if the access point is far away. The other concern is, of course, the interference between access points. The power adjustment is usually done through configuration and, in some cases, through a monitoring function on the network. In the latter case, the monitoring function reports the information to a central controller.

A new initiative within the IEEE (802.11k) has been started to improve traffic distribution within the network. Specifically, it addresses the problem of access point overloading so that stations can connect to underused access points for a more efficient use of network resources.

With respect to power management on the client side, a station can indicate that it is going into a “sleep” or low-power state to the access point through a status bit in a frame header (refer to Figure 1b). The access point then buffers packets for the station instead of forwarding them to the station as soon as they are received. The sleeping station periodically wakes up to receive beacons from the access point. The beacons include information about whether frames are being buffered for the station. The station then sends a request to the access point to send the buffered frames. After receiving the frames, the station can go back to sleep.

802.11a/g Technology—Orthogonal Frequency-Division Multiplexing

Sometimes called *discrete multitone* (DMT) in the *Digital Subscriber Line* (DSL) world, *Orthogonal Frequency-Division Multiplexing* (OFDM) is used as the underlying technology in 802.11g and 802.11a. OFDM is a form of *Frequency-Division Multiplexing* (FDM); normally, FDM uses multiple frequency channels to carry the information of different users. OFDM uses multicarrier communications, but only between one pair of users—that is, a single transmitter and a single receiver.

Multicarrier communications splits a signal into multiple signals and modulates each of the signals over its own frequency carrier, and then combines multiple frequency carriers through FDM. OFDM uses an approach whereby the carriers are totally independent of (orthogonal to) each other. Note that the total bandwidth consumed with OFDM is the same as with single carrier systems even though multiple carriers are used—because the original signal is split into multiple signals. OFDM is more effective at handling narrowband interference and problems related to multipath fading, simplifying the building of receiver systems.

We can illustrate this process with a simple example—one often used in discussions about OFDM. For a “normal” transmission at 1 Mbps, each bit can take 1 microsecond to send. Consider bit 1 and bit 2 sent with a gap of 1 microsecond. If two copies of bit 1 are received at the destination, one of them is the reflected or delayed copy. If the delay is around 1 microsecond, this delayed copy of bit 1 can interfere with bit 2 as it is received at the destination because they arrive at approximately the same time. Now consider an OFDM transmission rate of 100 kbps, that is, the bits are sent “slower” but over multiple frequencies. A multipath delay of around 1 microsecond will not affect bit 2, because bit 2 is now arriving much slower (around 10 microseconds). The delay in bit arrival (1 microsecond in our example) is not a function of the transmission—rather it is due to the various paths taken by the signal.

Orthogonal Frequency-Division Multiple Access (OFDMA) superimposes the multiple-access mechanism on OFDM channels, so that multiple users can be supported through subsets of the subcarriers assigned to different users. Note that 802.16-2004 (“Fixed” WiMAX) uses OFDM, whereas 802.16e-2005 (“Mobile” WiMAX) uses OFDMA.

MIMO and 802.11n

Multiple Input Multiple Output (MIMO) antennas are the basis for the 802.11n wireless LAN standard, currently in draft form but on the way to final standardization. Signals often reflect off objects and are received at different times and strengths at the receiver, resulting in a phenomenon called *multipath distortion*. (Note: 802.11n in this article implies the draft 802.11n standard at the time of writing.) MIMO actually takes advantage of this distortion by sending a single data stream split into multiple parts to be transmitted from multiple antennas (typically 3 in 802.11n) and letting the reflected signals be processed at the receiver (through multiple antennas). The transmission of multiple data streams over different spatial channels, sometimes known as *Space Division Multiplexing* (SDM), also allows a larger amount of data to be sent over the air. Through advances in the *Digital Signal Processing* (DSP)-based processing, the receiver can process the signals, cross-correlate them, and reconstitute them accurately despite interference. Also, because of the multiple signals received over multiple paths, link reliability is increased.

The 802.11n standard uses three antennas and also supports two radios (for the 2.4- and 5-GHz bands where 802.11n can operate). It can also use 40-MHz channels through *channel bonding*—that is, two adjacent 20-MHz channels are combined into a single 40-MHz channel, possibly resulting in a data rate of up to 150 Mbps of effective throughput.

One concern with 802.11n that is starting to gain attention is the power requirement of 802.11n access points. With radios in both bands and the use of MIMO, 802.11n access points tend to consume more power than the 802.11 a/b/g access points, leading to problems when the access point is powered by *Power over Ethernet (PoE)* power-sourcing equipment. The 802.3af standard permits a maximum of 12.95W per Ethernet port, which is often less than the power that most 802.11n APs need. The IEEE 802.3at working group is working toward a higher-power PoE standard. This initiative, commonly called *PoE Plus*, will peak at 25W per Ethernet port (on Category 5 Ethernet cable).

Ethernet Backhaul

The access point has two primary functions—connecting wireless clients to each other as well as connecting wireless and wired clients. In the latter, the access point can act as an Ethernet bridge by passing Layer 2 frames between the wired and wireless networks, or as a router, terminating WLAN and Ethernet Layer 2 frames and performing IP-level forwarding. The Layer 3 routing model is less popular and we will not consider it here.

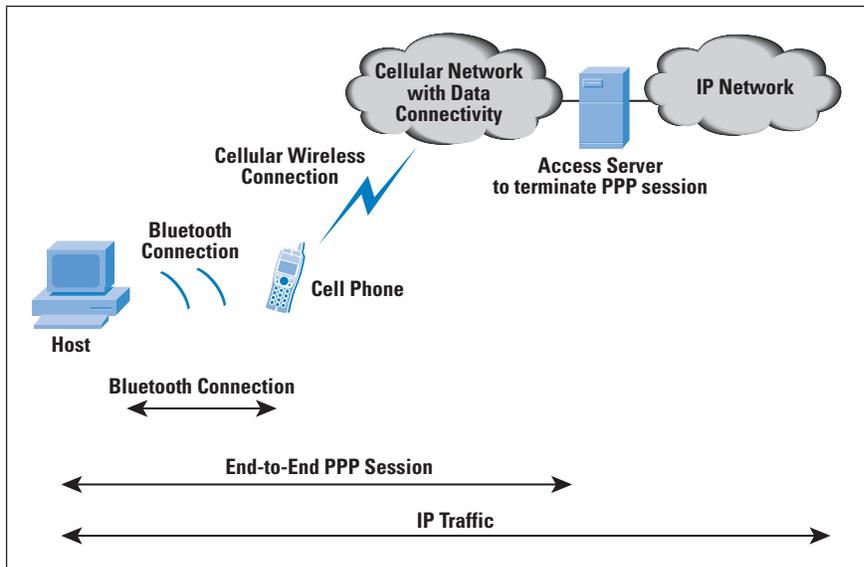
The access point typically terminates WLAN management and control frames. However, there is another model of a *thin access point* wherein these frames can be backhauled to a WLAN switch for processing. The access point connection to the wired network is typically an Ethernet link to a dedicated Ethernet switch port at 100-Mbps or Gigabit Ethernet speeds. With the advent of 802.11g and 802.11a WLANs, 10-Mbps links are not sufficient because these WLANs can operate at close to 27-Mbps throughput over the wireless network.

When considering 802.11n, we find that 100-Mbps backhaul links to the switch are insufficient for the 802.11n throughput of 150, or even 300 Mbps with channel bonding. Gigabit Ethernet links are often considered for connectivity between the 802.11n access point and the Ethernet switch. The next speed for Ethernet connectivity is 10 Gbps, which is well-established in the enterprise for data center and core Ethernet network applications. Work is ongoing in the IEEE for 40- and 100-Gbps Ethernet, so that should cover advances in wireless speeds for efficient backhaul to the wired network.

Bluetooth

Bluetooth started as a “wire-replacement” protocol for operation at short distances. A typical example is the connection of a phone to a PC, which, in turn, uses the phone as a modem (see Figure 2). The technology operates in the unlicensed 2.4-GHz ISM band. The standard uses FHSS technology. There are 79 hops in BT displaced by 1 MHz, starting at 2.402 GHz and ending at 2.480 GHz.

Figure 2: Typical Use of a Bluetooth enabled phone as a data modem for a PC



Bluetooth belongs to a category of *Short-Range Wireless* (SRW) technologies originally intended to replace the cables connecting portable and fixed electronic devices. It is typically used in mobile phones, cordless handsets, and hands-free headsets (though it is not limited to these applications). The specifications detail operation in three different power classes—for distances of 100 meters (long range), 10 meters (ordinary range), and 10 cm (short range).

Bluetooth operates in the unlicensed ISM band at 2.4 GHz (similar to 802.11 b/g wireless), but it is most efficient at short distances and in noisy frequency environments. It uses FHSS technology—that is, it avoids interference from other signals by hopping to a new frequency after transmitting and receiving a packet. Specifically, 79 hops are displaced by 1 MHz, starting at 2.402 GHz and finishing at 2.480 GHz.

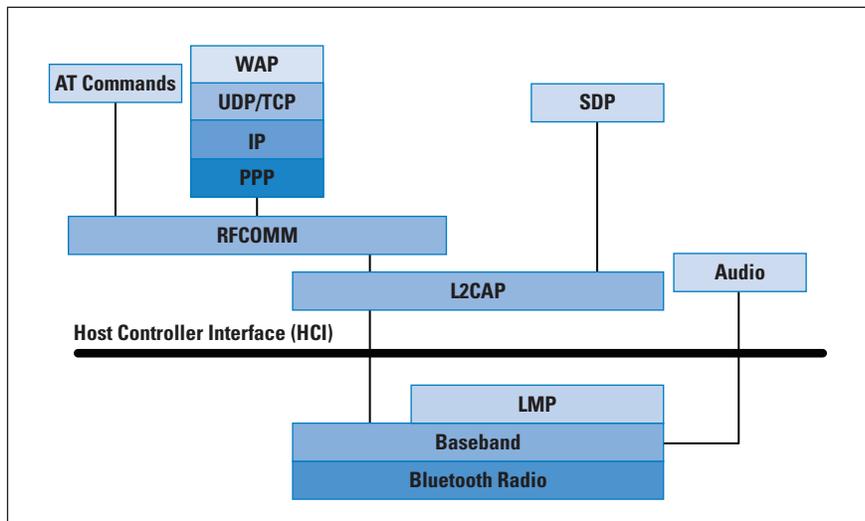
Bluetooth can operate in both point-to-point and logical point-to-multipoint modes. Devices using the same BT channel are part of a *piconet* that includes one master and one or more slaves. The master BT address determines the frequency hopping sequence of the slaves. The channel is also divided into time slots, each 625 microseconds in duration. The master starts its transmission in even-numbered time slots, whereas the slave starts its transmission in odd-numbered slots.

BT specifies two types of links, a *Synchronous Connection-Oriented* (SCO) link and an *Asynchronous Connectionless Link* (ACL). The SCO link is a symmetric point-to-point link between a master and a single slave in the piconet, whereas the ACL link is a point-to-multi-point link between the master and all the slaves participating in the piconet. Only a single ACL link can exist in the piconet, as compared to several individual SCO links.

Bluetooth Stack

Other than the radio and baseband components (the physical layer of Bluetooth that manages physical channels and links), the Bluetooth stack (see Figure 3) includes a *Link Manager Protocol* (LMP) used for link management between the endpoints, a *Logical Link Control and Adaptation Protocol* (L2CAP) for the data link, a *Radio Frequency Communication* (RFCOMM) protocol to provide emulation of serial ports over L2CAP, and a *Service Discovery Protocol* (SDP) for the dynamic discovery of services—because the set of services changes dynamically based on the RF proximity of the devices. In addition, the *Host Controller Interface* (HCI) provides a uniform command interface to the baseband controller and the link manager to have access to the hardware registers.

Figure 3: Key Elements of the Bluetooth Stack



LMP is required for authentication, encryption, switching of roles between master and slave, power control, and so on. L2CAP provides both connection-oriented and connectionless data services functions, including protocol multiplexing, segmentation and reassembly, and piconet-based group abstraction. As part of the multiplexing function, L2CAP uses the concept of channels, with a channel ID representing a logical channel endpoint on a BT device. L2CAP offers services to the higher layers for connection setup, disconnect, data reading and writing, pinging the endpoint, and so on.

RFCOMM, which provides emulation of serial ports on the BT link, can support up to 60 simultaneous connections between two BT devices. The most common emulation is of the RS-232 interface, which includes emulation of the various signals of this interface such as *Request To Send* (RTS), *Clear To Send* (CTS), *Data Terminal Ready* (DTR), and so on. RFCOMM is used with two types of BT devices—endpoints such as printers and computers and intermediate devices such as modems. In Figure 3, the IP stack over *Point-to-Point Protocol* (PPP) over RFCOMM emulates the mode of operation over a dialup or dedicated serial link. Because the various BT devices in a piconet may offer or require a different set of services, the *Service Discovery Protocol* (SDP) is used to determine the nature of the services available on the other nodes. SDP uses a request-response packet scheme for its operation.

Bluetooth Profiles

BT includes multiple profiles that correlate to the type of services that are available from BT nodes. For example, the BT headset profile is used between an audio source and a headset, both connecting wirelessly through BT—it involves a subset of the well-known AT commands used with modems. The audio source (typically a cell phone or cordless phone) implements the BT audio gateway profile for communicating with the device implementing the headset profile. Other profiles include a basic printing profile (often used for printing between a PC and a BT-enabled printer), dialup networking profile, fax profile, cordless telephony profile, *Human Interface Device* (HID) profile, and so on. The last profile is used for BT-enabled keyboards and mice—it is based on the HID protocol defined for USB.

The Bluetooth dialup networking profile is interesting from an IP perspective; as shown in Figures 2 and 3, it involves the IP stack running over RFCOMM to provide the appearance of a serial port running PPP, which is very similar to dialup networking over a basic telephone service line.

Bluetooth Frame Format and Speeds

The frame format in BT consists of a 72-bit field for the access code (including a 4-bit preamble, 64-bit synchronization field, and 4 bits of trailer), followed by a 54-bit header field that includes information about the frame type, flow control, acknowledgement indication, sequence number, and header error check. Following the header field is the actual payload, which can be up to 2745 bits. In all, the frame length can be a maximum of 2871 bits. Whereas synchronous BT traffic has periodic reserved slots, asynchronous traffic can be carried on the other slots.

BT ranges can vary from a low-power range of 1 meter (1 mW) for Class 3 devices, 10 meters (2.5 mW) for Class 2 devices, to 100 meters (100 mW) for Class 1 devices. BT Version 1.2 offers a data rate of 1 Mbps, and BT Version 2.0 with *Enhanced Data Rate* (EDR) supports a data rate of 3 Mbps. BT Version 1.1 was ratified as the IEEE Standard 802.15.1 in 2002.

Bluetooth versus Wi-Fi

A few years ago, some marketing literature tried to emphasize BT and Wi-Fi as competing technologies. Though both operate in the ISM spectrum, they were invented for different reasons. Whereas Wi-Fi was often seen as a “wireless Ethernet,” BT was initially seen purely as a cable- or wire-replacement technology. Uses such as dialup networking and wireless headsets fit right into this usage model. Recently, the discussion has focused more on coexistence instead of competition because they serve primarily different purposes. There are still some concerns related to their coexistence because they operate over the same 2.4-GHz ISM band.

To recapitulate, the Bluetooth physical layer uses FHSS with a 1-MHz-wide channel at 1600 hops/second (that is, 625 microseconds in every frequency channel). Bluetooth uses 79 different channels. Standard 802.11b/g uses DSSS with 20-MHz-wide channels—it can use any of the 11 20-MHz-wide channels across the allocated 83.5 MHz of the 2.4-GHz frequency band. Interference can occur either when the Wi-Fi receiver senses a BT signal at the same time that a Wi-Fi signal is being sent to it (this happens when the BT signal is within the 22-MHz-wide Wi-Fi channel) or when the BT receiver senses a Wi-Fi signal.

BT 1.2 has made some enhancements to enable coexistence, including *Adaptive Frequency Hopping* (AFH) and optimizations such as Extended SCO channels for voice transmission within BT. With AFH, a BT device can indicate to the other devices in its piconet about the noisy channels to avoid. Wi-Fi optimization includes techniques such as dynamic channel selection to skip those channels that BT transmitters are using. Access points skip these channels by determining which channels to operate over based on the signal strength of the interferers in the band. Adaptive fragmentation is another technique that is often used to aid optimization. Here, the level of fragmentation of the data packets is increased or reduced in the presence of interference. For example, in a noisy environment, the size of the fragment can be reduced to reduce the probability of interference.

Another way to implement coexistence is through intelligent transmit power control. If the two communicating (802.11 or Wi-Fi) devices are close to each other, they can reduce the transmit power, thus lowering the probability of interference with other transmitters.

WiBree to Low-Energy Bluetooth

WiBree is a technology first proposed by Nokia to enable low power communication over the 2.4-GHz band for button cell (or equivalent) battery-powered devices. A consequence of the low power requirement is the need for the wireless function to perform a very small set of operations when active and go back to the sleep or to standby mode when inactive.

The WiBree technology has been adapted by the *Bluetooth Special Interest Group* (SIG) as part of the lower-power BT initiative—also known as *Low Energy* (LE) BT technology. The LE standard is expected to be finalized sometime in 2009. When this standardization is completed, three types of BT devices will be available: traditional BT, LE BT, and a mixed or dual-mode BT. A mixed-mode device can operate in low power mode when communicating with other LE devices (for example, sensors) and traditional BT mode when communicating with BT devices, implying the presence of both a BT stack and an LE stack on the same device.

WiMAX

WiMAX stands for *Worldwide Interoperability for Microwave Access* and is defined under the IEEE 802.16 working group. Two standards exist for WiMAX—802.16d-2004 for fixed access, and 802.16e-2005 for mobile stations^[9]. The WiMAX forum certifies systems for compatibility under these two standards and also defines network architecture for implementing WiMAX-based networks.

WiMAX can be classified as a last-mile access technology similar to DSL, with a typical range of 3 to 10 kilometers and speeds of up to 5 Mbps per user with non-line of sight coverage. WiMAX access networks can operate over licensed or unlicensed spectra in various regions or countries—though licensed spectrum implementations are more common. WiMAX operation is defined over frequencies between 2 and 66 GHz, parts of which may be unlicensed spectrum deployments in some countries. The lower frequencies can operate over longer ranges and penetrate obstacles, so initial network roll-outs are in this part of the spectrum—with 2.3-, 2.5-, and 3.5-GHz frequency bands being common. Channel sizes vary from 3.5, 5, 7, and 10 MHz for 802.16d-2004 and 5, 8.75, and 10 MHz for 802.16e-2005. WiMAX networks are often used to backhaul data from Wi-Fi access points. In fact, they are often envisaged as replacements for the current implementation of metro Wi-Fi networks that use 802.11b/g for client access and 802.11a for backhaul to connect to the other parts of the network.

Technology

The 802.16d-2004 standard uses OFDM similar to 802.16a and 802.16g, whereas 802.16e-2005 uses a technology called *Scalable Orthogonal Frequency Division Multiplexed Access* (S-OFDMA). This technology is more suited to mobile systems because it uses subcarriers that enable the mobile nodes to concentrate the power on the subcarriers with the best propagation characteristics (because a mobile environment has more dynamic variables). Likewise, the 802.16e radio and signal processing is more complex.

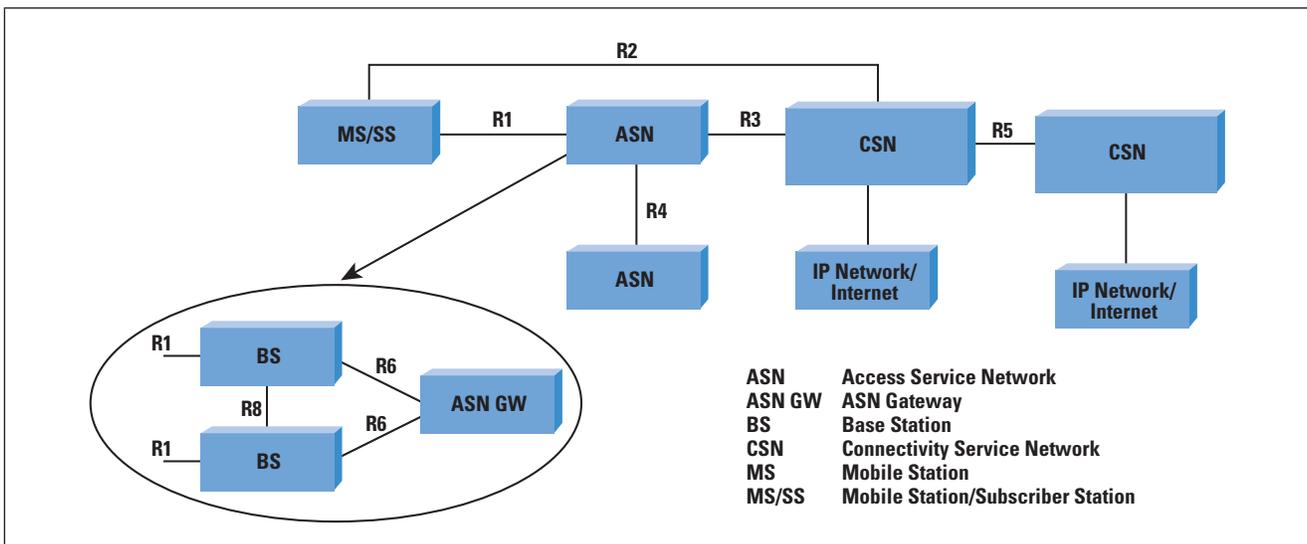
Unlike 802.11, which supports only *Time-Division Duplexing* (TDD)—where transmit and receive functions occur on the same channel but at different times, 802.16 offers TDD, *Frequency-Division Duplexing* (FDD) (transmit and receive on different frequencies, which could also be at different times). Another innovation in WiMAX is similar to the scheme in *Code Division Multiple Access* (CDMA)—subscriber stations are able to adjust their power based on the distance from the base station, unlike the case of client stations in an 802.11 network.

WiMAX base stations use a scheduling algorithm for medium access by the subscriber stations. This access is through an access slot that can be enlarged or contracted (to more or fewer slots) that is assigned to the subscriber stations. *Quality-of-Service* (QoS) parameters can be controlled through balance of the time-slot assignments among the base stations. The base-station scheduling types can be unsolicited grant service, real-time polling service, non-real time polling service, and best effort. Depending upon the time of traffic and service requested, one of these scheduling types can be used.

WiMAX Network Architecture

The WiMAX network architecture is specified through functional entities (see Figure 4), so you can combine more than one functional entity to reside on a network element. The *Mobile Station* (MS) connects the *Access Service Network* (ASN) through the R1 interface—which is based on 802.16d/e. The ASN is composed of one or more *base stations* (BSs) with one or more ASN gateways to connect to other ASNs and to the *Connectivity Service Network* (CSN). The CSN provides IP connectivity for WiMAX subscribers and performs functions such as *Authentication, Authorization, and Accounting* (AAA)^[10,11], ASN-CSN tunneling, inter-CSN tunneling for roaming stations, and so on. A critical tenet of the WiMAX Forum network architecture is that the CSN must be independent of the protocols related to the radio protocols of 802.16.

Figure 4: WiMAX Forum Network Architecture Functional Blocks and Interface Points



The R3 interface (reference point) is used for the control-plane protocols and bearer traffic between the ASN and CSN for authentication, policy enforcement, and mobility management. The base station connects to an ASN gateway to provide the MS with external network access. The R6 interface between the BS and ASN-GW could be open or closed based on the profile—in fact, you could have a co-located base station and *ASN gateway* (ASN-GW), depending upon the network implementation. The ASN gateway uses the R3 interface to communicate with the AAA services in the visited CSN (that is, the CSN “corresponding” to the ASN). The servers in the visited CSN can communicate with the home CSN (that is, the CSN corresponding to the “home” network of the MS). In the simplest case multiple ASNs (WiMAX networks) connect through ASN gateways to the public Internet (that is, there is only one *Network Service Provider* (NSP) and the visited and home CSNs are the same). Note that you could implement a WiMAX network with just one ASN and one CSN—in that case, the R3 interface would be completely internal and not exposed.

Three profiles are identified to map ASN functions into ASN-GW and BS functions. These profiles are considered an implementation guideline for how you would build the various devices implementing these functions. Profile A is a strict separation of the BS and ASN-GW functions, where the ASN-GW controls and manages radio resources that are located on the BS and also provides the handover and data-path functions. The R6 interface is exposed in this profile.

Profile B is a more integrated function, where the BS has more functions than in profile A; in fact, the BS might even integrate most of the ASN functions. The R6 interface is a closed interface in this profile. The third profile is profile C, which is similar to profile A except that the base stations incorporate more functions, including radio resource management and control as well as hand-offs.

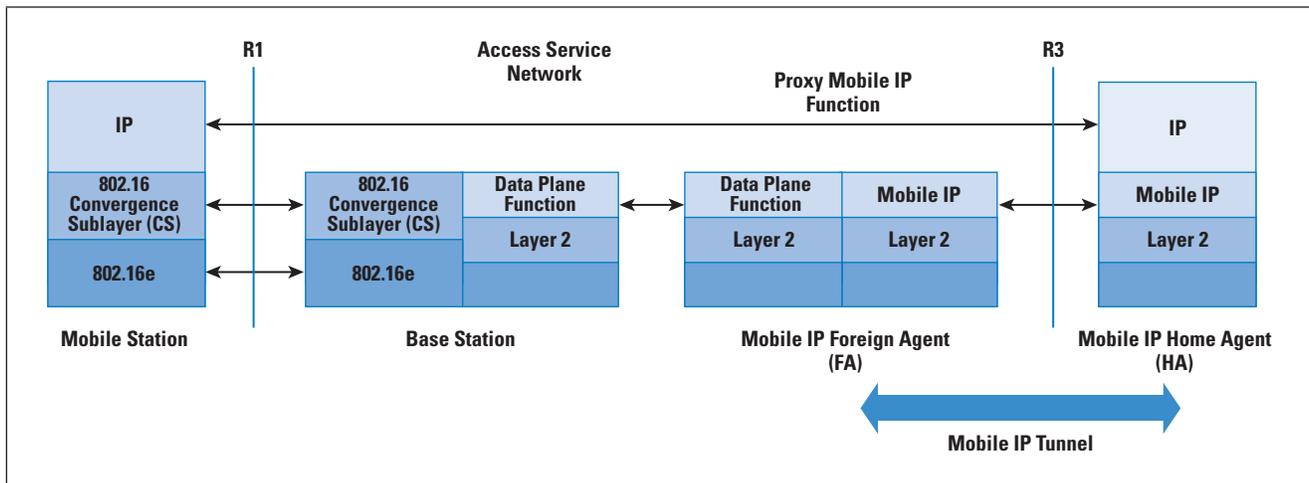
IP Connectivity and Data Transfer

The MS can be a fixed IP gateway (think of an 802.11 access point that provides connectivity to users in a coffee shop and connects to the IP network of the service provider through WiMAX) or a mobile end node (for example, a laptop with WiMAX connectivity). The IP address used by the gateway on the connection to the WiMAX network is known as the *Point of Attachment* (PoA) address. A third type of access is nomadic access, where the IP gateway can be moved from one location to another but connects to the network only after it has been relocated.

When the station is mobile, the WiMAX Forum specifies that the *Mobile IP* (MIP) architecture and protocols should be used. There are two types of Mobile IP possible: *Client Mobile IP* (CMIP) and *Proxy Mobile IP* (PMIP). The former involves changes to the MS protocol stack, but the latter does not.

The architecture can support both models. In the P-MIP scenario (see Figure 5), the ASN implements the Foreign Agent (see William Stallings' article in IPJ on Mobile IP^[8]), and terminates Mobile IP tunnels for the various mobile stations in the same ASN.

Figure 5: Data Transport and Proxy Mobile IP in WiMAX



In the figure, the MS has an address at the point of attachment that is used to forward packets from the MIP Foreign Agent inside the ASN. Because the ASN acts as a proxy of the attached MS, this implementation is known as a *Proxy MIP* implementation—also, there is no need for the MS to be aware of the MIP function being performed by the network.

Perspective on WiMAX versus Cellular Services

The WiMAX Forum has specified that the *Network Working Group* (NWG) architecture should be capable of supporting voice, multimedia services, and priority services such as emergency voice calls. It also supports interfacing with interworking and media gateways. Also, the service permits more than one voice session per subscriber, as well as simultaneous voice and data sessions. Support of IP Broadcast and Multicast services over WiMAX networks is also included. The architecture is also expected to support differentiated QoS levels at a per-MS or -user level (coarse grained) and at a per-service flow (fine-grained) level. It shall also support admission control and bandwidth management.

Initially, WiMAX was touted by some as a replacement for cellular services. An important consideration was using *Voice over IP* (VoIP) for voice calls—that is, where voice was another service over the data network. This model was in contrast to the existing cellular service where data was an adjunct to the basic service of TDM-based voice. More recently, WiMAX is being positioned as a data-connectivity option for remote locations, especially where it would be difficult to lay new copper or optical cable. Not surprisingly, these options are being pursued aggressively in developing countries.

Common Misperceptions About Wi-Fi, BT, and WiMAX Technologies

We have considered the key aspects of the three technologies—Wi-Fi, BT, and WiMAX—and their position in IP networks. In this section, we will outline and clarify some common perceptions and misperceptions about these technologies.

1. *BT and Wi-Fi are competing technologies*—Actually, they address a different set of requirements despite operating in the same 2.4-GHz space. BT is a “wire replacement” usually for short distances. Wi-Fi is typically used for data, voice, and video traffic over distances up to 300 meters.
2. *WiMAX is Wi-Fi on steroids*—To clarify, this statement is an oversimplification used often in the trade press. WiMAX operates in licensed spectra and uses a different network architecture as compared to Wi-Fi, which is in the unlicensed spectrum and uses a simple access point to wired Ethernet architecture. One overlapping function is for backhauling Wi-Fi traffic, which can be done by Wi-Fi (typically 802.11a) or WiMAX.
3. *Unlike BT, Wi-Fi cannot be used for voice*—This perception is not true because you can send multimedia traffic over Wi-Fi networks implementing 802.11e QoS functions that rely on the access point and stations implementing priority-based traffic transmission and scheduling.
4. *Wireless networks are not secure*—Although there is some validity to this argument because it is easier to eavesdrop on wireless networks, implementation of security schemes such as *Wi-Fi Protected Access* (WPA/WPA2) will help alleviate this problem.
5. *Wireless and radio technologies consume more power*—This statement is often true if the devices transmit continuously or have to increase their power because of the distance between the transmitter and receiver. Noisy channels contribute to this power use also. However, with careful engineering of the wireless implementation and techniques such as power save (in Wi-Fi) and short duty cycle transmissions, the power requirement can be lowered.

Summary

In this article, we have provided a flavor for IEEE 802.11 WLAN, Bluetooth, and WiMAX technologies and their implementation—specifically, how the nodes on these networks connect to an IP network. These technologies often serve complementary functions for end-to-end connectivity.

For Further Reading

- [1] IEEE 802.11 Standard, <http://standards.ieee.org/get-ieee802/download/802.11-2007.pdf>
- [2] Edgar Danielyan, “IEEE 802.11,” *The Internet Protocol Journal*, Volume 5, No. 1, March 2002.
- [3] T. Sridhar, “Wireless LAN Switches—Functions and Deployment,” *The Internet Protocol Journal*, Volume 9, No. 3, September 2006.
- [4] IEEE 802.16-2004 IEEE Standard for Local and metropolitan area networks Part 16: Air Interface for Fixed Broadband Wireless Access Systems.
- [5] IEEE 802.163-2005 IEEE Standard for Local and metropolitan area networks Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems Amendment for Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands, IEEE, <http://standards.ieee.org/getieee802/802.16.html>
- [6] Bluetooth Special Interest Group Publications, <http://www.bluetooth.com/Bluetooth/Technology/>
- [7] http://www.wimaxforum.org/technology/documents/WiMAX_Forum_Network_Architecture_Stage_2-3_Rel_1v1.2.zip
- [8] William Stallings “Mobile IP,” *The Internet Protocol Journal*, Volume 4, No. 2, June 2001.
- [9] Jarno Pinola and Kostas Pentikousis “Mobile WiMAX,” *The Internet Protocol Journal*, Volume 11, No. 2, June 2008.
- [10] Convery, S., “Network Authentication, Authorization, and Accounting – Part One: Concepts, Elements, and Approaches,” *The Internet Protocol Journal*, Volume 10, No. 1, March 2007.
- [11] Convery, S., “Network Authentication, Authorization, and Accounting – Part Two: Protocols, Applications, and the Future of AAA,” *The Internet Protocol Journal*, Volume 10, No. 2, June 2007.

T. SRIDHAR is Vice President of Technology at Flextronics in San Jose, California. He received his BE in Electronics and Communications Engineering from the College of Engineering, Guindy, Anna University, Madras, India, and his Master of Science in Electrical and Computer Engineering from the University of Texas at Austin. He can be reached at T.Sridhar@flextronics.com

The End of Eternity

Part One: IPv4 Address Exhaustion and Consequences

by Niall Murphy, Google, and David Wilson, HEAnet

“Eternity is a very long time, especially towards the end,” said Woody Allen^[22,23], and he was mostly right. The eternity that the 32 bits of IPv4 address space promised is now almost at an end, and we are faced with the task of deciding what to do after the “end of eternity.”

The size of the problem of IPv4 exhaustion is, unfortunately, also proportional to its longevity^[1,2,3]. Although the next-generation (IPng) effort^[4] kick-started the development of IPv6 partially in response to concern about the IPv4 consumption rate, the industry as a whole largely ignored the problem after *Classless Inter-Domain Routing* (CIDR) and the *Regional Internet Registries* (RIR) system contained the depletion problem to a manageable horizon. More recently, after Geoff Huston’s^[5] work showing that the expected depletion time was sooner than many organizations had expected, the concern has received considerable attention in address-allocation policy circles.

In this article, we examine IPv4 exhaustion in more detail. We talk about what exactly exhaustion will mean and what we can do about it, and then present a vision for the postexhausted world. Those familiar with our RIPE-55 talk^[6] will find much that is familiar, but the arguments have been expanded for a more general audience. The authors, as in that talk, are speaking only for themselves, and not their organizations.

What Does Exhaustion Mean?

Trivially, the point of IPv4 exhaustion is the point at which the guaranteed-free-and-unused pool runs out and the current allocation mechanism comes to an end. Although the depletion of the free pool defines the technical point of exhaustion, it is not the depletion itself that is of primary importance. After all, if it were, we could simply declare a moratorium on allocations with immediate effect, to preserve the resource for some notional future requirements. Rather, it is the effect on the practices and procedures, within the RIRs and within the *Local Internet Registries* (LIRs), administrative and technical, that will practically define exhaustion. These practices, which have grown to fit around the current behavior of the addressing system, the free pool, and so on, will require urgent reform after exhaustion, as indeed will the RIR system in general.

Currently organizations use and require new addresses for essentially every IP-related additional deployment (for example, adding customers to a publicly numbered DSL service, adding extra *Secure Sockets Layer* (SSL)-enabled websites to a Web hosting service, and adding extra publicly reachable servers to almost any service).

It has been emphasized that this problem affects only the *growth* of organizations performing IP deployments^[7]. Although it is important to acknowledge the partial correctness of this statement, much about the postexhaustion state could undermine the stability of well-established advertisements and routes unless the transition is well-handled. It seems intuitively correct that those who received allocations before exhaustion will be unaffected by exhaustion turmoil^[8], but we regard this premise as optimistic, as you will see later.

Along those lines, one less well-examined consequence of exhaustion is the erosion of the consensus model of Internet governance. There is potential for wide divisions to open up at the local and regional level unless this consensus is carefully conserved. No clear successor to the current model as yet exists; the RIRs appeared to be heading toward a spectrum of positions on, for example, the allocation of the last portions of the IPv4 free pool^[9, 10] until quite recently^[24].

The erosion of this model of governance as a consequence of exhaustion has been neither widely examined nor expected in the Internet community. Partially, this situation arose because of the useful and well-executed role that the RIRs have historically filled in providing sensible and stable conditions for decision making; some proportion of the membership of the RIRs might well feel that IPv4 exhaustion is a problem like any other, which the RIRs themselves are in the perfect position to resolve. However, although the atmosphere of mutual cooperation fostered by the RIRs has produced many useful service-related outputs (for example, the *Test Traffic Measurement* service of *Réseaux IP Européens* [RIPE]^[25]), one of the major non-obvious benefits they have brought is to provide a centralized focus for discussion with governments and regulatory agencies. Not only is it more efficient and therefore less time-wasting to centralize through one representative organization, it has also created expectations that similar matters can be dealt with in the same coordinated way—a very valuable expectation, which has helped to increase the credibility of industry self-regulation. This credibility allowed, for example, the *Number Resource Organization* (NRO) to help forestall a proposal to allocate IPv6 according to geographical boundaries^[27, 28].

Indeed, without credible industry self-regulation, it is not at all clear that this community could have grown as fast as it did. Although it seems clear today that the RIRs are the correct place for this kind of activity to go on (witness RIPE's "enhanced cooperation" task force^[26]), if they had not been around, government would either have had to deal with an organization with less of a pedigree or one with more inherent bias, or multiple organizations with competing biases, all of which could compel them to distrust the results of their liaisons. Unfortunately, in this respect the RIRs have been a victim of their own success.

Just as the consensus model in domains broke down when top- and second-level domains became monetized, so it is likely that the inherent win or loss for any given holder in any policy changes will undermine attempts to build consensus for address policy in a monetized IPv4 world. Absent this consensus, many of the RIR services that we rely upon will be undermined—not least the veracity of the WHOIS database and subsequent reliability of our routing filters, but also the RIR and *Internet Corporation for Assigned Names and Numbers* (ICANN) representations toward governments.

What Are the Problems with Exhaustion?

The biggest problem is the simplest one: existing organizations whose business model or operations are *solely* predicated on an ongoing flow of IPv4 addresses will fail. This premise would seem an extreme, even theoretical, characterization, but the size of this category in the real world is larger than you might think. Numerous organizations are also in trouble, perhaps less predicated upon IPv4 than the others, but that—for example—might have financial or operational difficulty in making the postexhaustion transition happen internally. They would also be placed at risk. Finally, there are those organizations that might rely on others to perform their transition correctly in order for them to continue effective operations: less directly at risk, but still probably affected.

Those who deal with the operation of the Internet on a daily basis are well-aware of the workarounds available that could save organizations from the doomsday scenario. It is unfortunate, then, that many of us have looked to the simplest cases in our immediate experience in order to form our opinions of the scale of the problem. It is indeed true that, in the short term, the client-side problem has largely been solved—provided that your customers or developers never have expectations in line with an end-to-end Internet. (It would seem that address-space pressure is likely to erode whatever end-to-end expectations still remain in today's Internet.)

However, the server-side problem (for example, SSL Website hosting, *IP Security* [IPsec] VPN endpoints, ...) remains unsolved. Workarounds exist^[11], but whether they will be ready and deployed in time remains an open question. There are, therefore, organizations operating at this moment that depend upon the continued availability of IPv4 addresses. Adequate workarounds have yet to be developed—never mind proven—for these businesses.

The situation becomes more complicated when we consider the candidate solutions. For example, such organizations as described previously cannot solve their problem by deploying IPv6 alone prior to the end of the transition, because they require universal reachability. Without universal reachability, support costs will rise, the quality of the user experience will decrease, and the credibility of Internet governance will be threatened. The only available evidence shows our position on the IPv6 transition curve being at the very beginning^[12].

Therefore it is difficult to emphasize this enough—new entrants providing Internet services *cannot expect to compete equally with existing operations*—because they have a very high barrier to entry formed not by the natural action and development of competitors, but by the resource scarcity of new addresses. Without new addresses, they cannot have an IPv4 *Default-Free Zone* (DFZ) routing-table entry; without a DFZ entry, they cannot be multihomed; without multihoming, they cannot offer sufficiently redundant Internet service; and without sufficiently redundant Internet service, they cannot meaningfully compete with existing operators.

A variety of poor-quality “fudges” are possible, of course: they could use the address space of their upstream operators (and run the risk of having that address space pulled or charged for), or they could outsource any address-requiring services to another organization (and be unable to control their service quality, as well as dependent upon their continuing operation), or they could host through some kind of public proxy network that redirects to their back-end servers through various hard-coded means (and create a fragile, difficult-to-operate network with higher running costs per unit customer than their competitors).

We will examine the other negative consequences of exhaustion in more detail later in the discussion; meanwhile, let us assume that the scenario described previously is undesirable enough for us to ask whether we can actually do anything to forestall it.

Can We Practically Defer Exhaustion?

What we would ideally like is some policy or algorithm that would give us more time—how much time is open to question—without producing its own set of ill effects. (We can certainly defer exhaustion by ceasing to allocate new IPv4 addresses tomorrow, but that solution is hardly practical.) Unfortunately, this problem is very difficult to resolve. Such direct precedents that appear clearly related to the current situation provide no useful guidance. Many resource-exhaustion problems have been faced before, but ultimately the solutions for those can be categorized into three kinds:

- *Make the resource renewable*: In this case, the resource is in danger of running out, but can be replenished by some means. Often this replenishment involves constraining production predicated on the resource to some smaller value, particularly when there is a natural rate of renewal—for example, fishing stocks. In the case of IPv4, it is fundamentally nonrenewable in that the resource is of a finite size. (As we discussed previously, current reclamation efforts^[13], although worthy of pursuit as a low-overhead task, cannot be a solution.)
- *Move to another resource*: This solution is already under way in the sense that we are engaged in the transition to IPv6. However, adoption of IPv6 will not happen fast enough to prevent the negative consequences of exhaustion.

- *Divide the resource more fairly:* This solution is useful primarily in the case where hoarding is taking place, causing resource problems for some significant proportion of a resource-using population. We are dividing the resource fairly as it is, and certainly since the emergence of the RIRs. For reasons discussed later, husbanding the resource more carefully is unlikely to actually be a solution.

We have faced other abstract exhaustion problems before as well: for example, phone-number depletion is somewhat similar to our current problem. However, phone-number depletion admits of a simpler solution—the creation of extra digits in the number space—because of the centralization of network knowledge in a comparatively small number of switches. For the Internet, where every deployed host would have to be informed about changes to the number space, such an approach is not operationally feasible. Furthermore, adding extra digits to the number code is not in fact simple, and telecommunications companies have experienced a wide range of problems with such approaches in the past, to say nothing of the loss of revenue and the failure of calls to connect because of customer confusion^[14,15]. We see no historical situation that provides a clear precedent and a clear way forward.

SimLIR

Accordingly, to help answer the question posed in the preceding section, we wrote a tool, *SimLIR*, to explore exhaustion and post-exhaustion scenarios. Rather than being a tool influenced primarily by computations based on growth curves, a “top-down” approach, it is a modeling tool that examines how changes in behavior affect relative consumption rates. Roughly 6,000 lines of Python, the tool is due to be open-sourced at its Google Code page^[16] shortly after this article is available. The tool models the whole *Internet Assigned Numbers Authority* (IANA)—>RIR—>LIR hierarchy, and currently maps LIRs to countries; it uses the same publicly available data as Geoff’s work. We would appeal to the community to help improve the program, because more research is desperately needed in this area.

Running the tool under various scenarios has produced preliminary results indicating that we cannot meaningfully defer exhaustion, given our current growth rates. It can be used to compare the effect of policy adjustments on known historical and simulated behavior. For example, one simple policy adjustment that has been informally suggested is to decrease the initial allocation size for new LIRs. Modeling this allocation with the tool, we halve the size the LIRs receive at the time of initial membership. If we allow this scenario to run to completion, we have seen that it allows us to defer exhaustion by less than a week. Intuitively, we might expect this assumption to be realistic because startup activity, although important, is relatively small in terms of proportion of allocations. New LIRs numbered approximately 500 in 2006^[17], and any scheme that attempted to defer exhaustion based on such a small proportion of overall operations could not practically succeed.

The question then arises whether any other scheme based upon treating some partition of the request-space differently could have a significant positive effect. However, such a scheme necessarily assumes that some set of requests are oversized, and can in fact be shrunk with no ill effects. Even if they are oversized, identifying them without inducing either unworkable bureaucracy or a chilling effect on the operations of the organization would be a significant task, not lightly undertaken. Furthermore, it would be in the self-interest of the current RIR membership not to agree to such a change in policy. With any such scheme, there would be a non-zero chance of their own requests being deemed faulty in some respect, thus leading to significant risk to their own operations. All of this process would of course be happening in the approach to exhaustion, where it would be more critical than ever to receive enough numbering resources! We can assume, therefore, that no such scheme would ever make it past the policy-making apparatus of bottom-up-influenced RIRs. Ironically, the easiest changes to enact are changes governing allocations to startup organizations; the affected organizations are not in the room at the time of policy formation, because they are not members yet. But such changes are highly unlikely to have a positive effect.

Finally, partitioning schemes are similar to other schemes proposed to rework the *End Game* for IPv4 allocation^[18, 19, 20] or retain a certain proportion of the free pool for as-yet-unknown future needs, in that we put RIRs in the awkward situation of having to decide that some requests are more legitimate than others, at a time when these requests are likely to be particularly urgent. RIRs should not be in the business of deciding who gets to have new customers, and partitioning the request space invites the possibility of preferential treatment. We can be sure that any preferential treatment at this crucial time, accidental or otherwise, would attract lawsuits. Judicial involvement in the allocation process close to the time of exhaustion would benefit almost nobody.

It is important to note that these risks are mainly specific to partitioning the request space from the RIR to the LIR; in other words, imposing criteria at the time of request. Partitioning the remaining pool per RIR, that is, imposing criteria at the time of division, such as proposed by the $n = 1$ policy^[24], does not suffer from “favoritism.” Indeed, even if there were blatantly iniquitous division at the IANA-to-RIR level, although various checks and balances exist to ensure there is not, it would be unlikely to affect those with resources sufficient to possess an office in the region in question, or to open one up; it is patently clear that the requests will follow where the space is, and it is highly unlikely that any single RIR with a large amount of space left after others have been exhausted would be in any kind of position to pass a discriminatory policy.

We make these points to highlight that any scheme based upon LIR partitioning presents immense difficulties of principle. Even if these difficulties are worked out, they seem unlikely to meaningfully defer exhaustion: the current run rate for IPv4 address space will exhaust the space within a 5-year timeframe anyway, even if all practically possible measures are taken.

The Consequences of Scarcity

Suppose for the moment that at the time of exhaustion, Internet-connected organizations have to fend for themselves, with no particularly well-defined industry strategy in place. We would then expect to see a broad movement within the industry to conserve precious public IPv4 address space. One obvious way for an organization to obtain more usable IPv4 space is to move previously publicly-numbered resources behind *Network Address Translation* (NAT) gateways. Other, less-legitimate sources of new addresses will probably also be explored, and these actions, combined with the generally uncoordinated changes, may well trigger the following negative consequences:

- *Inability to measure clients, and difficulty of supporting them:* As we see more layers of NAT within networks, it becomes gradually more difficult to establish who is actually connecting to you, and what problems they are having. Cookies are a partial solution for only one important protocol. Measurement becoming harder means that support costs will rise.
- *Address-space hijacking:* As organizations become more desperate for space, it is entirely feasible that they will begin to cast around for space not explicitly unavailable in order to meet their business needs. How widespread this practice would be remains an open question, but effective barriers to this behavior are not currently available. We would expect a general deterioration in the quality of routing.
- *WHOIS database quality down:* Coupled with layers of NAT hiding more and more networks from direct sight, transfers of address space (legitimate or otherwise) will cause the WHOIS database to become gradually less and less accurate, leading to...
- *Distributed denial-of-service (DDoS) tracking trouble:* Problems tracking DDoS attacks and abuse origins of all kinds make law enforcement and network operators equally unhappy.
- *Connection quality down:* Connection quality, in terms of connections that complete successfully and have tolerable latency, will go down as a function of client growth behind gateways.
- *RIR billing model under pressure:* The RIRs will need to find a new way to pay their costs or go out of business—gradually, but inevitably. Of course the RIRs, like every other organization, must serve a need, but they currently provide a large number of ancillary services not directly related to IP allocation, and those services would also be under threat.

- *Consensus undermined*: This consequence is possibly the most dangerous of them all. If a chaotic state of affairs is allowed to continue for too long, our very ability to make decisions as a community will be undermined as organizations abandon the RIR model that has failed them. We will have squandered, in a way, the foundation of trust that allows such ethical codes as we have developed in Internet operations to persist. That foundation will not be easily recovered.

(Note that all of these are effects that are likely to emerge to varying degrees with the onset of scarcity, however it takes place; in other words, if the RIRs engage in a program of scarcity management by partitioning requests, it is highly likely that the scenario described previously will happen no matter what is left in the free pool.)

In any large shock such as we describe, there will be operational turmoil. Organizations will attempt to employ the technologies they need to dig themselves out of trouble, or bend the rules to the same end. There will be financial turmoil as the ability of each business to scale in the new regime is tested. Turmoil for existing businesses and new entrants will no doubt attract increased attention from governmental and quasigovernmental agencies of all kinds. Turnover in the routing table will increase as uncoordinated deaggregation of prefixes takes place. Unwelcome as all these consequences are, we will probably be far too preoccupied with our own individual problems to take care of the broader picture.

Postexhaustion Vision

Although we hope it is clear, given the previous discussion—that IPv4 addresses will still be required after exhaustion—our highest aspiration cannot be an Internet confined in perpetuity to IPv4 alone. If we are to continue in a manner resembling our current operations, we require continued address plenty, even by today’s rather restricted standards. The End Game, therefore, is an IPv6 Internet, or at least enough of one to keep off address scarcity for a workable subset of the industry.

So, the problem can then be characterized as the transition toward this state of affairs—the gap between the end of the old allocation model and the emergence of an adequate replacement. Any solution will have to either make the gap shorter, by bringing users to the IPv6 Internet sooner, or make it less painful, by helping IPv4-dependent organizations survive. (Note that a solution that makes the gap less painful may well cause it to lengthen.)

With the problem stated this way, we can evaluate possible solutions in this context. A hurried, stimulated transition of popular services to IPv6 will quite likely shorten the gap, although a mass transition is also likely to be an unstable one and so rather painful.

A voluntary release of unused addresses may help reduce the pain, but is unlikely to service the run rate adequately, given its voluntary nature, and in any event will prolong dependence on IPv4, thus lengthening the gap. Tweaking policies to make remaining IPv4 addresses arbitrarily difficult to get merely introduces the effects of scarcity still sooner, helping neither goal.

That said, our initial examination of the problems of exhaustion indicate that there will be a group of people who will require IPv4 addresses after the exhaustion point, and it is also clear that there are those who have addresses, such as the lucky recipients of class A addresses in the early days, but no particular incentive to give them up. We do not actually want to recycle these prefixes indefinitely, however; that just sustains the current model. Optimally, we should provide whatever opportunity we can to those who require IPv4 addresses, to get them (and us) toward the End Game of an adequate global IPv6 deployment.

We do not require an unlimited IPv4 supply to accomplish this goal. We do, however, require liquidity: the ability to transfer, with incentives to transfer. Although it is very difficult for a centralized system (such as an RIR) to reclaim adequate space, the effort/reward ratio is much more favorable for an individual organization that knows its own network. So we must provide some stimulus for them to increase liquidity, while imposing some realistic restriction on demand. It must of course be scrupulously fair.

Stated in this way, a market-based trading exchange is not just one way of attempting to solve the problem—such an exchange, properly regulated, is arguably the most neutral and fairest way to manage the problem of scarcity.

In the next article we will explore how such a market system should work, discuss what new problems it is likely to create, and consider the potential effect on the routing table.

References

- [1] <ftp://ftp.ietf.org/ietf-online-proceedings/94dec/area.and.wg.reports/ipng/ale/ale-minutes-94dec.txt>
- [2] <http://tools.ietf.org/html/rfc2008>
- [3] Hain, Tony, "A Pragmatic Report on IPv4 Address Space Consumption," *The Internet Protocol Journal*, Volume 8, No. 3, September 2005
- [4] <http://playground.sun.com/ipv6/doc/history.html>
- [5] <http://ipv4.potaroo.net>
- [6] <http://www.ripe.net/ripe/meetings/ripe-55/presentations/murphy-simlir.pdf>
- [7] http://www.isoc.org/educpillar/resources/ipv6_faq.shtml
- [8] <http://www.ietf.org/internet-drafts/draft-narten-ipv6-statement-00.txt>
- [9] <http://www.apnic.net/meetings/24/program/sigs/policy/presentations/el-nakhal-prop-051.pdf>
- [10] <http://www.ripe.net/ripe/policies/proposals/2007-06.html>
- [11] http://www.switch.ch/pki/meetings/2007-01/name-based_ssl_virtualhosts.pdf
- [12] For example, http://h.root-servers.org/128.63.2.53_2.html versus http://h.root-servers.org/h2_5.html
- [13] <http://www.ripe.net/ripe/meetings/ripe-55/presentations/vegoda-reclaiming-our.pdf>
- [14] A "smooth and convenient" dialing plan for India.
<http://www.mycoordinates.org/indias-phone-june-06>
- [15] http://en.wikipedia.org/wiki/UK_telephone_code_misconceptions
- [16] <http://code.google.com/p/simlir/>
- [17] <http://www.ripe.net/docs/ripe-407.html#membership>
- [18] <http://www.ripe.net/ripe/policies/proposals/2007-03.html>
- [19] <http://www.ripe.net/ripe/policies/proposals/2007-06.html>

- [20] <http://www.ripe.net/ripe/policies/proposals/2007-07.html>
- [21] <http://kuznets.fas.harvard.edu/~aroth/alroth.html>
- [22] Woody Allen, "Side Effects," 1980.
- [23] Woody Allen through (most famously) Stephen Hawking, <http://www.cnn.com/2006/WORLD/asiapcf/07/04/talkasia.hawking.script/index.html>
- [24] <http://icann.org/en/announcements/proposal-ipv4-report-29nov07.htm>
- [25] <http://www.ripe.net/ttm/>
- [26] <http://www.ripe.net/ripe/tf/enhanced-cooperation/index.html>
- [27] <http://www.nro.net/documents/nro18.html>
- [28] <http://www.ripe.net/maillists/ncc-archives/im-support/2004/index.html>
- [29] Huston, G., "The Changing Foundation of the Internet: Confronting IPv4 Address Exhaustion," *The Internet Protocol Journal*, Volume 11, No. 3, September 2008.

NIALL MURPHY holds a B.Sc. in Computer Science and Mathematics from University College Dublin. While in university, he founded the UCD Internet Society, which provided Internet access to approximately 5,000 students. He went on to work for (and found) various organizations: the **.IE** domain registry, Club Internet (now Magnet Entertainment), Ireland On-Line, Enigma Consulting, Bitbuzz, and Amazon.com. He is currently in Site Reliability Engineering at Google. He is the coauthor of numerous articles, some RFCs, the O'Reilly book *IPv6 Network Administration*, and is a published poet and keen amateur landscape photographer. E-mail: niallm@avernus.net

DAVE WILSON holds a B.Sc. in Computer Science from University College Dublin, not coincidentally from around the same time as Niall. He has worked at HEAnet, the Irish National Research & Education Network, for more than 10 years, maintaining an involvement with RIPE and with the pan-European research network Géant. Dave is a member of the ICANN Address Supporting Organization Address Council; he helped to found the Irish IPv6 task force, which has the support of the national government there. E-mail: dave.wilson@heanet.ie

Remembering Jon: Looking Beyond the Decade

by Vint Cerf, Google

A decade has passed since Jon Postel left us.^[0] It seems timely to look back beyond that decade and to look forward beyond a decade hence. It seems ironic that a man who took special joy in natural surroundings, who hiked the Muir Trail and spent precious time in the high Sierras, was also deeply involved in that most artificial of enterprises, the Internet. As the *Internet Assigned Numbers Authority* (IANA)^[1] and the *Request for Comments* (RFC) editor, Jon could hardly have chosen more polar interests. Perhaps the business of the artificial world was precisely what stimulated his interest in the natural one.

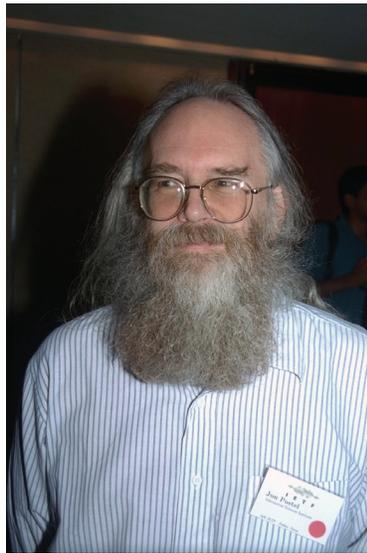


Photo: Peter Löthberg

As a graduate student at UCLA in the late 1960s, Jon was deeply involved in the ARPANET project, becoming the first custodian of the RFC note series inaugurated by Stephen D. Crocker. He also undertook to serve as the “Numbers Czar,” tracking domain names, Internet addresses, and all the parameters, numeric and otherwise, that were critical to the successful functioning of the burgeoning ARPANET and, later, Internet protocols. His career took him to the east and west coasts of the United States but ultimately led him to the University of Southern California’s *Information Sciences Institute* (ISI), where he joined his colleagues, Danny Cohen, Joyce K. Reynolds, Daniel Lynch, Paul Mockapetris, and Robert Braden, among many others, who were themselves to play important roles in the evolution of the Internet.

It was at ISI that Jon served longest and as the end of the 20th century approached, began to fashion an institutional home for the work he had so passionately and effectively carried out in support of the Internet. In consultation with many colleagues, but particularly with Joseph Sims of the Jones Day law firm and Ira Magaziner, then at the Clinton administration White House, Jon worked to design an institution to assume the IANA responsibilities. Although the path to its creation was rocky, the *Internet Corporation for Assigned Names and Numbers* (ICANN)^[2] was officially created in early October 1998, just two weeks before Jon’s untimely death on October 16.

In 1998 an estimated 30 million computers and 70 million users were on the Internet. In the ensuing decade, the user population has grown to almost 1.5 billion and the number of servers on the Internet now exceeds 500 million (not counting episodically connected laptops, *personal digital assistants* [PDAs], and other such devices). As this decade comes to a close, the *Domain Name System* (DNS) is undergoing a major change to accommodate the use of non-Latin character sets in recognition that the world's languages are not exclusively expressible in one script^[7]. A tidal wave of newly Internet-enabled devices as well as the increasing penetration of Internet access in the world's population is consuming what remains of the current IPv4 address space, accelerating the need to adopt the much larger IPv6 address space in parallel with the older one. More than three billion mobile devices are in use, roughly 15 percent of which are already Internet-enabled.

Jon would take considerable satisfaction knowing that the institution he worked hard to create has survived and contributed materially to the stability of the Internet. Not only has ICANN managed to meet the serious demands of Internet growth and importance in all aspects of society, but it has become a worked example of a new kind of international body that embraces and perhaps even defines a multi-stakeholder model of policy making. Governments, civil society, the private sector, and the technical community are accommodated in the ICANN policy development process. By no means a perfect and frictionless process, it nonetheless has managed to take decisions and adapt to the changing demands and new business developments rooted in the spread of the Internet around the globe.

Always a strong believer in the open and bottom-up style of the Internet, Jon would also be pleased to see that the management of the Internet address space has become regionalized and that five *Regional Internet Registries* (RIRs)^[3] now cooperate on global policy, serving and adapting to regional needs as they evolve. He would be equally relieved to find that the loose collaboration of DNS root zone operators has withstood the test of time and the demands of a much larger Internet, showing that their commitment has served the Internet community well. Jon put this strong belief into practice as he founded and served as ex-officio trustee of the *American Registry for Internet Numbers* (ARIN)^[4].

As the first individual member of the Internet Society he helped to found in 1992, Jon would certainly be pleased that it has become a primary contributor to the support of the Internet protocol standards process, as intended. The Internet Architecture Board and Internet Engineering and Research Task Forces, as well as the RFC editing functions, all receive substantial support from the Internet Society.

He might be surprised and pleased to discover that much of this support is derived from the Internet Society's creation of the *Public Interest Registry* (PIR)^{5, 6} to operate the `.org` top-level domain registry. The Internet Society's scope has increased significantly as a consequence of this stable support, and it contributes to global education and training about the Internet as well as to the broad policy developments needed for effective use of this new communication infrastructure.

As a computer scientist and naturalist, Jon would also be fascinated and excited by the development of an interplanetary extension of the Internet to support manned and robotic exploration of the Solar System. In October 2008, the Jet Propulsion Laboratory began testing of an interplanetary protocol using the Deep Impact spacecraft now in eccentric orbit around the sun. This project began almost exactly 10 years ago and is reaching a major milestone as the first decade of the 21st century comes to an end.

It is probable that Jon would not agree with all the various choices and decisions that have been made regarding the Internet in the last 10 years, and it is worth remembering his philosophical view: "Be conservative in what you send and liberal in what you receive."

Of course he meant this idea in the context of detailed protocols, but it also serves as a reminder that in a multi-stakeholder world, accommodation and understanding can go a long way toward reaching consensus or, failing that, at least toleration of choices that might not be at the top of everyone's list.

No one, not even someone of Jon's vision, can predict where the Internet will be decades hence. It is certain, however, that it will evolve and that this evolution will come, in large measure, from its users. Virtually all the most interesting new applications of the Internet have come not from the providers of various Internet-based services, but from ordinary users with extraordinary ideas and the skills to experiment. That they are able to experiment is a consequence of the largely open and nondiscriminatory access to the Internet that has prevailed over the past decade. Maintaining this spirit of open access is the key to further development, and it seems a reasonable speculation that if Jon were still with us, he would be in the forefront of the Internet community in vocal and articulate support of that view.

A 10-year toast seems in order. Here's to Jonathan B. Postel, a man who went about his work diligently and humbly, who served all who wished to partake of the Internet and to contribute to it, and who did so asking nothing in return but the satisfaction of a job well done and a world open to new ideas.

References

- [0] Vint Cerf, “I Remember IANA,” *The Internet Protocol Journal*, Volume 1, No. 3, December 1998. Also published as RFC 2468, October 1998.
- [1] <http://www.iana.org>
- [2] Vint Cerf, “Looking Toward the Future,” *The Internet Protocol Journal*, Volume 10, No. 4, December 2007.
- [3] Daniel Karrenberg, Gerard Ross, Paul Wilson, and Leslie Nobile, “Development of the Regional Internet Registry System,” *The Internet Protocol Journal*, Volume 4, No. 4, December 2001.
- [4] <http://www.arin.net>
- [5] <http://www.pir.org>
- [6] <http://www.isoc.org>
- [7] Huston, G., “Internationalizing the Domain Name System,” *The Internet Protocol Journal*, Volume 11, No. 1, March 2008.

VINTON G. CERF is vice president and chief Internet evangelist for Google. Cerf served as a senior vice president of MCI from 1994 through 2005. Widely known as one of the “Fathers of the Internet,” Cerf is the co-designer of the TCP/IP protocols and the architecture of the Internet. He received the U.S. National Medal of Technology in 1997 and the 2004 ACM Alan M. Turing award. In November 2005, he was awarded the Presidential Medal of Freedom. Cerf served as chairman of the board of the *Internet Corporation for Assigned Names and Numbers* (ICANN) from 2000 through 2007 and was founding president of the Internet Society. He is a Fellow of the IEEE, ACM, the American Association for the Advancement of Science, the American Academy of Arts and Sciences, the International Engineering Consortium, the Computer History Museum, and the National Academy of Engineering. He is an honorary Freeman of the City of London. Cerf holds a Bachelor of Science degree in Mathematics from Stanford University and Master of Science and Ph.D. degrees in Computer Science from UCLA. E-mail: vint@google.com

Letters to the Editor

IPv4 Address Exhaustion

I read with interest your article in *The Internet Protocol Journal* (Volume 11, No. 3, September 2008) regarding the IPv4 address exhaustion problem. It occurs to me that two approaches for encouraging the public and *Internet Service Provider* (ISP) community to migrate to IPv6 are being dismissed somewhat, but used creatively together might offer some hope for pushing us in that direction: government regulation and changing the fact that there isn't a public interest in IPv6.

What if government regulation forced a new or currently existing common service to use IPv6? One obvious possibility is video content. Since the broadcast industry is already regulated by the FCC, further regulation providing for governance of this type of application isn't too much of a stretch. Consumer demand is likely to increase in this area as broadband continues to be widely deployed, and if the public were required to run in dual-stack mode to access it, the likelihood of adoption would be much greater. It would also incent the ISPs to provide connectivity to the IPv6 address space, possibly even with a revenue-generating model behind it.

I reluctantly bring up the pornography industry as another type of content that could be relegated to the IPv6 address space. It is my understanding that this type of traffic as a percentage of the total is quite large. Based on this assumption, it would have the same effect of forcing the large portions of the public and ISPs to provide connectivity to the IPv6 address space. Again, I mention this industry reluctantly, but from a political perspective regulation of this industry and its content is likely to be an easier proposal for the public to support since you could use the "value" of disconnected portions of the Internet to best advantage.

I realize that the global nature of the Internet makes regulation and the subsequent enforcement extremely difficult. But, I also assume that even if our enforcement were controlled only at the perimeter of the U.S. traffic it would have a strong effect on the behavior of the public and ISPs.

Best regards,

—John Newell, INX Inc.
jcnnewell@gmail.com

The author responds:

Thanks for your response. It is true to say that various efforts have been undertaken across many years to find a "killer-app" for IPv6, if I may be permitted to use that overabused and by now very tired term. To date these efforts have not been successful. That's not because of any lack of trying.

There have been some really quite innovative ideas for IPv6 over the years, and so far most of them have been retrofitted into IPv4 one way or another. From one perspective this retrofit is entirely logical, given that good ideas tend to thrive in locations where audiences are receptive, and today's IPv4 Internet is still a very fertile place for good ideas to flourish.

The other part of the problem is that service providers tend to create innovative services with existing markets in minds, so these days the novel applications and services that appear to gain the attention of significant parts of the user base tend to operate in the IPv4 network, and by necessity such applications and services account for *Network Address Translation* (NAT) devices and various forms of filters and firewalls.

These observations indicate that a certain reinforcing cycle exists that cements the existing role of the IPv4 Internet, and tends to work against the widespread deployment of innovative services that are feasible only in the IPv6 environment.

So if the adoption of IPv6 is a carrot or stick affair, our efforts to find some tempting carrots have, so far, not been overly successful. We've been unable to identify particular goods or services for which there is a compelling case of consumer demand coupled with a set of technology constraints that imply that the service is feasible only across a deployed IPv6 infrastructure with IPv6 endpoints. So if the field we are working in is bereft of carrots, are there any available sticks that we can use instead? In this case there is the same old stick that originally motivated IPv6 in the first place: We are running out of IPv4 addresses. If we believe that there is more to do in the Internet, more people to connect, more devices to add, more conversations to have, more services to deploy, more ideas to realize, and more objectives to achieve, then IPv4 cannot in and of itself sustain that vision for the Internet. The threat here is that the growth of the IPv4 Internet may well cease when the supply of further IPv4 addresses is exhausted.

Is this threat of network stagnation going to be enough to propel us into an IPv6 Internet? Will it be an adequate motivator to encourage the necessary investment in network infrastructure and in the provision of goods and services that first operate in a transitional dual-stack environment, and ultimately in an IPv6 world? I hope that the answers are "yes," as do many others I'm sure.

But I'm also worried that it may not be enough and that we may spin off into an entirely different trajectory that ultimately dismantles most of the attributes of today's Internet. I worry that instead of an open network that fosters innovation and creativity we might end up with "vertical integration" and "transparent convergence" and a network that actively resists new services and applications.

So for me, and I hope many others, IPv6 needs no new “killer-app.” IPv6 does not need television or pornography to succeed. IPv6 is an imperative for the Internet simply because the alternatives to IPv6 appear to offer us a leap backward in technology and a leap backward in the elastic ways we’ve been able to use networks—and in the process we are going to destroy the Internet as we know it!

Regards,

—Geoff Huston, APNIC
gih@apnic.net

Dear Ole,

In his latest IPJ article (Volume 11, No. 3), Geoff Huston highlights the significance of NAT as a mechanism enabling service providers to externalize the costs and risks arising from IPv4 address scarcity. While acknowledging the increased burden and uncertainty borne by end users and NAT-traversing applications, Geoff speculates that the success of this mechanism is likely to inspire the deployment of yet another level of (“carrier grade”) address translation, to further prolong if not absolutely preclude the incorporation of IPv6 by incumbent service providers. While entirely plausible, such a move would create the same kind of “double blind” conditions for Internet service delivery that prevailed in financial markets when debt securitization was coupled with the externalization of asset depreciation risks in the form of *Credit Default Swaps*. In such cases, the second layer of indirection tends to make it all too easy to maintain self-serving assumptions (and/or plausible deniability) about the true nature and purpose of the first layer, and thus to fuel the perpetuation of unsustainable industry practices unto the point of industry collapse. Given the now inescapable lessons of the recent financial sector collapse, it would be nice if we didn’t have to learn this particular one again the hard way.

—Tom Vest
tvest@eyeconomics.com

On Paper

I just received the September issue (Volume 11, No. 3) of IPJ and wanted to make a quick comment about the paper change. Upon reading the section on the change I quickly dug up the previous copy of IPJ and compared the two. I personally like the new paper much better. The main reason I like it is because it is much easier on the eyes, I think mostly because it no longer has a glare from overhead lighting reflecting like the old paper type did. It’s a welcomed change from my take.

—David Swafford,
Network Engineer for CareSource, Dayton, OH, US
david@davidswafford.com

Book Reviews

A Dictionary and a Handbook

Hundreds of telecom books are published each year, but it is unusual to find a really good one. There must have been a blue moon (I'll have to check my almanac) this month, for I found two new and quite remarkable books by the same author, Ray Horak. One is a dictionary and the other an encyclopedic work, both covering the full range of voice, data, fax, video, and multimedia technologies and applications that comprise contemporary telecommunications. Further, they do so in such a plain-English, commonsense manner that you don't need to be a serious telecom student or professional to benefit from them—any layperson with a serious need to know will find them to be of great value. Finally (and this is rare in a technical book), both are actually relatively easy and certainly interesting reads, with liberal doses of fascinating historical context. In fact, they are even strong on entertainment value, with humorous observations and quotations sprinkled throughout. Horak has written each book in a different style for a different purpose, so they are best acquired together—as a set.

Webster's New World Telecom Dictionary

Webster's New World Telecom Dictionary, by Ray Horak, ISBN-10: 047177457X, ISBN-13 978-0471774570, Wiley Publishing Inc., 2007.

In order to communicate effectively in a contemporary telecom conversation, one must speak a special language rife with technical terminology, much of which is in the form of abbreviations, acronyms, contractions, initialisms and portmanteaux. To add to the confusion, many terms have multiple very precise—and occasionally imprecise—meanings, depending on the context. Writing a telecom dictionary must be a formidable task, one which only either the very brave or very foolhardy would even attempt. I'm not sure into which category Ray Horak falls, but his *Webster's New World Telecom Dictionary* is an excellent piece of work.

Organization

Dictionaries are in alphabetical order, of course, with chapters thrown in for symbols and numbers. Because the introduction of symbols requires special treatment, within each of the 28 chapters Horak organizes the approximately 4,600 definitions in ASCII order, perhaps as an accommodation for the binarians among us. The book includes an appendix of standards organizations and special interest groups, which can be useful if you need more information on a subject or need to know exactly to whom to complain about a *standard* or *specification*, both of which terms are defined clearly in the dictionary, of course.

Comparisons: Comprehensive and Correct

In my opinion, the best telecom dictionary ever written, aside from *Webster's*, is the *Communications Standard Dictionary*, by Martik H. Weik. That book unfortunately is out of print, with the final 3rd edition dated 1996. At 1095 pages, it is a bit overwritten and way too technical for most purposes, reading much like an IEEE dictionary. At this point, it certainly is out-of-date.

A handful of other telecom dictionaries and encyclopedias are currently in print, by far the most popular of which is *Newton's Telecom Dictionary*. Because *Newton's* dominates the market and has done so for many years, any telecom dictionary or encyclopedia is inevitably compared to that work. *Webster's New World Telecom Dictionary* is no exception, particularly because Ray Horak was the contributing editor to *Newton's* from the 12th through the 22nd editions.

Although *Webster's* defines only 4,600 terms in comparison to *Newton's* highly dubious claim of some 24,500 terms, *Webster's* definitions are much better researched, much more precise, and much more efficiently worded (that is, there is much less “fluff”). Even if *Webster's* almost certainly will gain in bulk as future editions expand the coverage of the telecom domain, it contains all of the essential telecom and IT terms, and defines them clearly and concisely. *Webster's* includes many humorous definitions but, unlike *Newton's*, they are all relevant and meaningful. For example, Horak lists three types of standards—*de jure*, *de facto*, and *du jour*. According to him, a *du jour* standard is defined as follows:

“From French, meaning *of the day*. The popular standard of the day. One day 10 years ago, ATM was really hot and a lot of people made a lot of money talking about ATM and selling products based on ATM. It seemed like only the next day that IP was really cool. (I made this one up.)”

Other humorous definitions include analogue, endianness, Hellenologophobia, hoot 'n' holler, OCD, PC, and WMBTOTCITB-WTNTALI. All of these, and more, serve to lighten the load, so to speak, but none of this humor detracts from what is a serious book on a serious subject. *Newton's*, on the other hand, is so full of personal observations and anecdotes, irrelevant humor (?), and inaccurate definitions as to make you wonder why bother to make the comparison at all. Horak states that he wrote *Webster's* partly to atone for his sins in contributing to *Newton's*, but mostly to put an authoritative reference book in his own hands, and those of others involved in litigation support. He apparently does a fair amount of work as an expert witness in intellectual property (the other IP) cases and on innumerable occasions has been asked to define and opine on terms such as link, circuit, channel, call, connection, switch, router, and PSTN. Now he can testify in court with one hand on the Good Book and the other on *Webster's*.

Recommended

Webster's New World Telecom Dictionary is an excellent piece of work. Ray Horak and his technical editor, Bill Flanagan, have collaborated to create a well-written, authoritative work that clearly sets a new standard for telecom dictionaries. I highly recommend it to anyone serious about telecom.

Telecommunications and Data Communications Handbook

Telecommunications and Data Communications Handbook, by Ray Horak, ISBN-10: 0470041412, ISBN-13: 978-0470041413, John Wiley & Sons, 2007.

Unless you have really big hands, you may wonder how it is that a tome of 791 pages that weighs more than 3 pounds could possibly be called a handbook. Well, the term “handbook” actually is fairly imprecise, but Ray Horak’s *Telecommunications and Data Communications Handbook* certainly is not. Actually, it is about as compact as it can be, given its encyclopedic nature, and it is very precise, indeed. The book covers the entire telecom landscape, from wireline to wireless, from copper to radio and fiber, from electrical to optical, and from the customer premises to the cloud. It discusses voice, data, fax, video and multimedia technologies, systems, and applications in great detail, and in the LAN, MAN, and WAN domains. The handbook explores every relevant technology, standard, and application in the telecom and datacom space.

Horak is a well-known telecom consultant, author, writer, columnist, and lecturer. The *Telecommunications and Data Communications Handbook* is based on his best-selling *Communications Systems and Networks* (1997, 2000, 2002), but is considerably more technical and broader in scope. It is exceptionally well-written in Horak’s plain-English, commonsense style, making it just as helpful to the neophyte and layperson as to the serious student or seasoned IT professional. Horak makes liberal use of well-constructed graphics to illustrate system and network architectures, topologies, and applications.

Organization

The Handbook begins with an excellent table of contents (20 pages) and ends with an excellent index (29 pages), both of which are crucial to a good book. After all, it doesn’t make any difference how good the information is if you can’t find it. The book is logically organized into 15 chapters and 2 appendixes.

Chapter 1 is devoted to fundamental concepts and definitions, thereby building a firm foundation of concepts and terminology upon which subsequent chapters build. Terms such as two-wire, four-wire, circuit, link, channel, switch, and router are clearly defined, compared, and contrasted. Chapter 2 explores the full range of transmission systems, including twisted pair (UTP, STP, and ScTP), coaxial, microwave, satellite, *Free Space Optics* (FSO), fiber-optics, *powerline carrier* (PLC), and hybrid systems.

Chapter 3 examines voice communications systems: KTS, PBX, Centrex, and ACD. Chapter 4 discusses messaging systems in detail, including facsimile (fax), voice processing, and e-mail and instant messaging, concluding with a detailed discussion of unified messaging and unified communications. Chapter 5 is dedicated to the *Public Switched Telephone Network* (PSTN) and addresses *Numbering Plan Administration* (NPA), regulatory domains, rates and tariffs, signaling and control systems, and network services. Chapter 6 returns to fundamentals, this time in the data communications domain, with detailed explanations of *Data Communications Equipment* (DCE) such as modems, codecs, CSUs, and DCUs, and then moves on to protocol basics, code sets, data formats, error control, compression techniques, network architectures, and security mechanisms.

Chapter 7 deals with conventional digital and data networks such as DDS, Switched 56, VPNs, T/E-carrier, X.25, and ISDN. Chapter 8 treats *Local-Area Networks* (LANs) and *Storage Area Networks* (SANs) exhaustively, including transmission media, topologies, broadband vs. baseband, equipment, operating systems, and standards. This chapter covers 802.3, 802.11, HiperLAN, Bluetooth, IEEE 1394, Fibre Channel, and iSCSI in considerable detail. Chapter 9 is devoted to broadband network infrastructure, including both access technologies (for example, xDSL, CATV, WLL, PON, and BPL) and transport technologies (for example, SONET/SDH and RPR). Chapter 10 offers an exhaustive study of broadband network services, including Frame Relay, ATM, Metropolitan Ethernet, B-ISDN, and AINs.

Chapter 11 discusses wireless, with an emphasis on mobility, covering both broad concepts and technical specifics of *Specialized Mobile Radio* (SMR), paging, cellular (1G, 2G, 2.5G, 3G, and beyond), packet data radio networks, and mobile satellite networks (GEOs, MEOs, and LEOs). Chapter 12 thoroughly treats video and multimedia networking, including a detailed discussion of video and multimedia standards (for example JPEG, MPEG, and H.320), *Session Initiation Protocol* (SIP), and IPTV. Chapter 13 exhaustively and insightfully explores the Internet and *World Wide Web* (WWW), including a thorough discussion of the IP protocol suite. Chapter 14 briefly examines convergence, and Chapter 15 examines telecom regulation, with a focus on the United States.

Appendix A is something of a decoder for abbreviations, acronyms, contractions, initialisms, and symbols. Appendix B gives a complete listing of relevant standards organizations and special interest groups, including full contact information, in case you need more information or want to offer comments on a particular subject.

Comparisons

It is hard to make a valid direct comparison to this book. *The Irwin Handbook of Telecommunications*, by James Harry Green, is good, but less complete, less technical, and drier, if such a combination is possible. The most recently published 5th edition also is apparently out of print. The *Voice & Data Communications Handbook*, by Regis “Bud” Bates, is written at a lower level; and, the *Essential Guide to Telecommunications*, by Annabel Dodd, at a much lower level. These latter two books are breezy reads and appeal more to a mass market than to a serious student or professional.

The *Telecommunications and Data Communications Handbook* compares more correctly to some of the more seminal works of Gilbert Held or James Martin, but covers a much wider range of subject matter and is a much easier and more pleasant read.

Recommended

The *Telecommunications and Data Communications Handbook* is written for the academic and professional community, but is just as relevant to anyone who needs to understand telecommunications system and network technologies and their meaningful applications. It is an exceptional work that should be on every IT professional’s bookshelf...when not in his or her hands.

—John R. Vacca,
jvacca@frognet.net

Read Any Good Books Lately?

Then why not share your thoughts with the readers of IPJ? We accept reviews of new titles, as well as some of the “networking classics.” In some cases, we may be able to get a publisher to send you a book for review if you don’t have access to it. Contact us at ipj@cisco.com for more information.

Fragments

Itojun Service Award Launched

A new award, providing recognition and support for those progressing IPv6 development on the Internet, was announced in November. The *Itojun Service Award* honors the memory of Dr. Jun-ichiro “Itojun” Hagino, who passed away in 2007, aged just 37^[1]. The award, established by the friends of Itojun and administered by the *Internet Society* (ISOC), recognizes and commemorates the extraordinary dedication exercised by Itojun over the course of IPv6 development. Itojun worked as a Senior Researcher at the *Internet Initiative Japan* (IIJ), was a member of the board of the *Widely Integrated Distributed Environment* (WIDE) Project, and from 1998 to 2006 served on the groundbreaking KAME project in Japan as the “IPv6 Samurai.” He was also a member of the *Internet Architecture Board* (IAB) from 2003 to 2005.

At the time of his passing, Russ Housley, *Internet Engineering Task Force* (IETF) Chair, and Olaf Kolkman, IAB Chair, issued a joint statement, praising Itojun’s service to IPv6 developments, saying that he had “inspired many and will be missed.”

The Itojun Service Award will run for 10 years, presented annually to an individual who has made outstanding contributions in service to the IPv6 community. The award includes a presentation crystal, a US\$3,000 honorarium, and a travel grant. The Award will honor an individual who has provided sustained and substantial technical contributions, service to the community, and leadership. With respect to leadership, the selection committee will place particular emphasis on candidates who have supported and enabled others in addition to their own specific actions.

The selection committee members for the Itojun Service Award are: Jun Murai, Hiroshi Esaki, Ole Jacobsen, Bob Hinden, Randy Bush, Bill Manning, Tatuya Jinmei, Kazu Yamamoto, and Kenjiro Cho.

Memorial donations to the Itojun Service Award Fund are welcomed and the Internet Society has established an account for donations. Details of the fund, as well as more information about Jun-ichiro “Itojun” Hagino and the Itojun Service Award are available on the ISOC Web site: <http://www.isoc.org/awards/itojun/>

The WIDE Project has also established a Japanese bank account to collect donations in Japanese Yen, the details of which are available here: <http://www.wide.ad.jp/itojun-award>

[1] Hinden, Bob, “Remembering Itojun: The IPv6 Samurai,” *The Internet Protocol Journal*, Volume 10, No. 4, December 2007.

EsLaRed Receives 10th Annual Postel Service Award

ISOC awarded the *Jonathan B. Postel Service Award* for 2008 to *La Fundación Escuela Latinoamericana de Redes* (EsLaRed) of Venezuela for its significant contributions to promote information technologies in Latin America and the Caribbean.

It is now ten years since the passing of Internet pioneer Jonathan B. Postel, the inspiration for this prestigious award. To mark this event in a special way, ISOC formed a *10th Anniversary Award Committee* including all the past award recipients, which has formally recognised EsLaRed for “its sustained efforts to bring scientific, technical, and social progress in Latin America and the Caribbean through education, research, and development activities on technology transfer.”

ISOC presented the award, including a US\$20,000 honorarium and a crystal engraved globe, in November during the 73th meeting of the IETF in Minneapolis, USA.

Accepting the award for EsLaRed was its President, Professor Ermanno Pietrosemoli. “We’re very excited to be honored in this way,” said Professor Pietrosemoli. “In the developing world, having access to the Internet, which gives us access to things like scientific journals and medical information, is not easy and it is not taken for granted. It is wonderful for us to be able to help people improve their conditions and to see first hand how the Internet can change people’s lives,” he said.

“On behalf of the ISOC community, it is my great pleasure to congratulate Professor Pietrosemoli and his dedicated colleagues at EsLaRed for their achievements over the years,” said Lynn St. Amour, President and CEO of ISOC. “EsLaRed’s commitment to the Internet has been at the forefront of regional development and their leadership has been an instrumental element in forming today’s dynamic Latin American and Caribbean Internet community,” said Ms St. Amour. For more information about this year’s recipient see:

<http://www.isoc.org/awards/postel/eslared.shtml>

The Postel Service Award was established by ISOC to honor individuals or organisations that, like Jon Postel, have made outstanding contributions in service to the data communications community. The award is focused on sustained and substantial technical contributions, service to the community, and leadership. Previous recipients of the Postel Award include Jon himself (posthumously and accepted by his mother), Scott Bradner, Daniel Karrenberg, Stephen Wolff, Peter Kirstein, Phill Gross, Jun Murai, Bob Braden and Joyce K. Reynolds (jointly), and Nii Quaynor. The award consists of an engraved crystal globe and a US\$20,000 honorarium. For more information see: <http://www.isoc.org/awards/postel/>

Call for Papers

The Internet Protocol Journal (IPJ) is published quarterly by Cisco Systems. The journal is not intended to promote any specific products or services, but rather is intended to serve as an informational and educational resource for engineering professionals involved in the design, development, and operation of public and private internets and intranets. The journal carries tutorial articles (“What is...?”), as well as implementation/operation articles (“How to...”). It provides readers with technology and standardization updates for all levels of the protocol stack and serves as a forum for discussion of all aspects of internetworking.

Topics include, but are not limited to:

- Access and infrastructure technologies such as: ISDN, Gigabit Ethernet, SONET, ATM, xDSL, cable, fiber optics, satellite, wireless, and dial systems
- Transport and interconnection functions such as: switching, routing, tunneling, protocol transition, multicast, and performance
- Network management, administration, and security issues, including: authentication, privacy, encryption, monitoring, firewalls, troubleshooting, and mapping
- Value-added systems and services such as: Virtual Private Networks, resource location, caching, client/server systems, distributed systems, network computing, and Quality of Service
- Application and end-user issues such as: e-mail, Web authoring, server technologies and systems, electronic commerce, and application management
- Legal, policy, and regulatory topics such as: copyright, content control, content liability, settlement charges, “modem tax,” and trademark disputes in the context of internetworking

In addition to feature-length articles, IPJ will contain standardization updates, overviews of leading and bleeding-edge technologies, book reviews, announcements, opinion columns, and letters to the Editor.

Cisco will pay a stipend of US\$1000 for published, feature-length articles. Author guidelines are available from Ole Jacobsen, the Editor and Publisher of IPJ, reachable via e-mail at ole@cisco.com

This publication is distributed on an “as-is” basis, without warranty of any kind either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. This publication could contain technical inaccuracies or typographical errors. Later issues may modify or update information provided in this issue. Neither the publisher nor any contributor shall have any liability to any person for any loss or damage caused directly or indirectly by the information contained herein.

The Internet Protocol Journal

Ole J. Jacobsen, Editor and Publisher

Editorial Advisory Board

Dr. Vint Cerf, VP and Chief Internet Evangelist
Google Inc, USA

Dr. Jon Crowcroft, Marconi Professor of Communications Systems
University of Cambridge, England

David Farber
Distinguished Career Professor of Computer Science and Public Policy
Carnegie Mellon University, USA

Peter Löthberg, Network Architect
Stupi AB, Sweden

Dr. Jun Murai, General Chair Person, WIDE Project
Vice-President, Keio University
Professor, Faculty of Environmental Information
Keio University, Japan

Dr. Deepinder Sidhu, Professor, Computer Science &
Electrical Engineering, University of Maryland, Baltimore County
Director, Maryland Center for Telecommunications Research, USA

Pindar Wong, Chairman and President
Verifi Limited, Hong Kong

*The Internet Protocol Journal is published quarterly by the Chief Technology Office, Cisco Systems, Inc. www.cisco.com
Tel: +1 408 526-4000
E-mail: ipj@cisco.com*

Copyright © 2008 Cisco Systems, Inc. All rights reserved. Cisco, the Cisco logo, and Cisco Systems are trademarks or registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. All other trademarks mentioned in this document or Website are the property of their respective owners.

Printed in the USA on recycled paper.



The Internet Protocol Journal, Cisco Systems
170 West Tasman Drive
San Jose, CA 95134-1706
USA

ADDRESS SERVICE REQUESTED

PRSR STD U.S. Postage PAID PERMIT No. 5187 SAN JOSE, CA
